

FERRY: access control and quota management service

Mine Altunay¹, Joseph Boyd¹, Bruno Coimbra¹, Kenneth Herner¹, Krysia Jacobs¹, Farrukh Kahn¹, Tanya Levshina^{1}, Brian McKittrick¹, Rennie Scott¹, Timothy Skirvin¹, Felix Stores¹, Jeny Teheran¹, Margaret Votava¹ and Tammy Whited¹*

¹Fermi National Accelerator Laboratory, Computing Sector, Batavia, IL, USA

Abstract. Fermilab developed the Frontier Experiments RegistRY (FERRY) service that provides a centralized repository for the access control and job management attributes such as batch and storage access policies, quotas, batch priorities and NIS attributes for cluster configuration. This paper describes FERRY architecture, deployment and integration with services that consume the stored information. The Grid community has developed several access control management services over the last decade. Over time, support for Fermilab experiments has required the collection and management of more access control and quota attributes. At the same time, various services used for this purpose, namely VOMS-Admin, GUMS and VULCAN, are being abandoned by the community. FERRY has multiple goals: maintaining a central repository for currently scattered information related to users' attributes, providing a Restful API that allows uniform data retrieval by services, and providing a replacement service for all the abandoned grid services. FERRY is integrated with the ServiceNow (SNOW) ticketing service and uses it as its user interface. In addition to the standard workflows for request approval and task creation, SNOW invokes orchestration that automates access to FERRY API. Our expectation is that FERRY will drastically improve user experience as well as decrease efforts spent on support by service administrators.

1 Introduction

The Fermilab Computing Sector (SC) provides comprehensive data processing and distributed computing framework for Fermilab's scientific stakeholders. The Frontier Experiments RegistRY (FERRY) service, developed by the SC, allows numerous services, including NIS, EOS, HTCondor, dCache and others, to have access to a single source of accurate information related to user identity mapping, authorization attributes and various quotas. Although it is not a core objective of the project, a side benefit is eliminating several obsolete services such as GUMS[1] and CMS LPC in-house developed service VULCAN[2] that provides fine-grained user mapping to specific resources. It also allowed us to retire the VOMS-ADMIN[3] tool that provides user interface to VOMS[4] service.

* Corresponding author: tlevshin@fnal.gov

The Fermilab Human Resource Database contains information on current and retired employees, visitors and summer students. Only active Fermilab users in good standing have a Fermilab Services account that provides access to Fermilab mail, SharePoint, and ticketing services. Only a subset of these users are people who need to run jobs on the Grid and transfer data to and from Fermilab storage services. This group of people is usually affiliated with one or several Fermilab experiments and projects (see Figure 1). Access to interactive and Grid clusters, to storage services and storage quotas, batch priority and slot allocation depends on their affiliation to a particular experiment. Prior to development and deployment of FERRY this information was been stored by various services and was often out of sync. FERRY created a single source of truth for this data.




Fig. 1. Groups of users that have been registered at FERRY, from Fermilab HR database to VO membership.

2 Requirements

Some of the essential FERRY service requirements, that came from the multiple stakeholders and customers, are listed below. FERRY service should:

- Provide flexible APIs that allow consuming services to pull information about:
 - VO role and grid map-files
 - VO members, their certificates, assigned groups and roles
 - LPC users, their CERN attributes, EOS quota and LPC group affiliations
 - Unix passwd and group files for a specified VO and/or compute resource
- Be able to extend and modify data schema. The schema should support:
 - users with multiple VO memberships and groups
 - resource-oriented groups that control access to a resource (NIS passwd and group files, storage access, ... etc.)
- Be available on 5x8 basis. The dependent services should rely on cached information.
- Provide appropriate level of security for data access. Though information that is stored by FERRY doesn't contain any sensitive information, it still gathers a lot of information related to users in one central place.
- Integrate with ServiceNow (SNOW) [5] to handle customer requests. Failure to propagate data from a customer request submitted via SNOW should be

periodically retried and a ticket to the Ferry support group should be opened when retry limit is reached.

- Must allow for removing a user or removing user attributes and affiliations.

3 Architecture

FERRY service consists of the RESTful Web Service and the FERRY database. FERRY Web Interface is SNOW. The SNOW custom forms have been created to fill specific requests. The SNOW service uses the “Orchestration”[6] technique to communicate with FERRY. Service providers custom scripts are querying the FERRY service using more than eighty APIs to get relevant information. Only small subset of clients are authorized to make modification via FERRY APIs. Clients are using x509 certificate to authenticate with FERRY Service.

Initial ingest of data was implemented by using several scripts that pull data from various sources, correlate and validate them, and then populate the FERRY database. The `update_user` script, run as a cron job, acquires information about newly joined or retired users from Fermilab’s HR database and Active Directory Service and then updates FERRY. All other modification requests come via SNOW. Fermilab services periodically query FERRY to get necessary data. The FERRY architecture is shown in Figure 2.

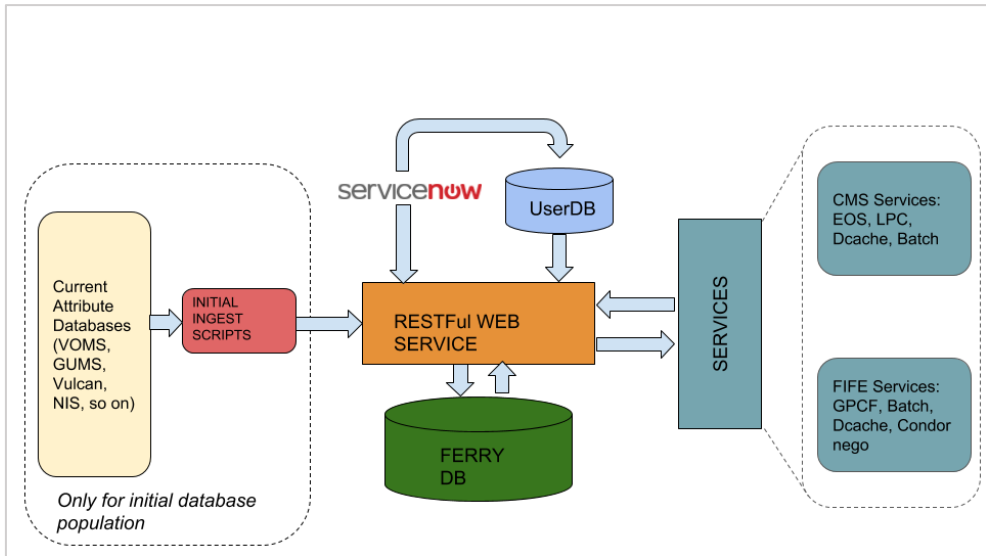


Fig. 2. FERRY Architecture.

4 Implementation

FERRY web service is implemented using GO language. Postgres was chosen for the database implementation. All FERRY APIs return information in JSON format. The databases tables are logically organized in groups:

- User related tables that contains user’s full name, unix id, unix name, and user’s groups.

- Affiliation unit (Virtual Organization), Fully Qualified Attribute Name (FQAN), members and certificates.
- Compute resource and access to it, that include user's primary and secondary groups, preferred shell and home directory.
- Storage resources, access to storage and quotas.

5 Integration with Services

The FERRY service substitutes existing authorization services, such as GUMS and VULCAN. This data, obtained from FERRY, allows correctly authorize users and provides valid mapping. FERRY enable us to get rid of stale and inaccurate data previously stored locally by various services.

Many grid services need a grid map-file that provides mapping of a distinguished name from a presented certificate to a local Unix account. This Unix account would be used for executing a job or accessing a file on a Unix node. Other services need a VO-role map-file for mapping a Fully Qualified Attribute Name to a local Unix account. The services that are in need of these files pull them from FERRY by using relevant FERRY APIs.

FERRY is keeping track of members affiliation with experiments. In the past, there were "multiple sources of truth" for this sort of information; for example NIS tables for experiment interactive clusters, HTCondor CE head and worker nodes all got information from different sources. FERRY also serves as a source of truth for populating VOMS instances. FERRY contains various other service configuration metadata which previously was stored only in configuration files of individual services. Since this data is now in FERRY, it can easily be modified by SNOW processes, initiated from user request forms, and the services can pull the configuration from the FERRY database at predetermined intervals to configure the underlying services. The information currently stored in FERRY includes HTCondor group account quotas and EOS storage quotas.

5.1 FERRY and Grid services

Fermilab provides computing and storage resources to multiple experiments. These resources are accessed through a gateway known as 'compute element' or CE. We use HTCondor-CE [7] as our gatekeeper of choice. Part of the responsibility of a CE service is determining if an experiment or a user is authorized to access the compute resources attached to it. HTCondor-CE at Fermilab relies on VO role and grid map-files to make this decision. VO role map-file contains VOMS attribute to Unix account mapping whereas grid mapfile contains X509 certificate distinguished name (DN) to Unix account mapping. FERRY, being the source of truth for access control, provides this information to the CE through RESTful API calls. A python script runs hourly on a caching server to fetch these map-files from FERRY and places them onto a local webserver. The CE service then periodically pulls in these map-files and uses them to make authorization decisions. We have ten CEs with different access control requirements. FERRY provides the necessary flexibility to enforce these access control policies per cluster.

5.2 FERRY and NIS

In order to tie FERRY to existing NIS services we use a caching server. This caching server pulls data from FERRY on an hourly basis; validates the data; converts the data into

"standard" data formats; and places the output onto a local web server. Each client can then query the data as necessary for its operations.

One example: user and group management replaces NIS with nss_db by using a JSON formatting compatible with json-passwd[8]. Each individual client (worker node, interactive machine, etc) independently pulls down group/passwd data for a specific resource from the FERRY caching server on an hourly basis, and saves it locally in a format compatible with the libnss_db Unix interface.

The major benefits of this architecture are:

- Simplicity - everything works through simple web calls and basic bash scripting.
- Performance - easily scales to thousands of systems.
- Reliability - clients can continue to get data during FERRY system outages and off-hours, and can use their existing cached data even with no further network access.

5.3 FERRY and VOMS

The synchronization script pulls relevant information from FERRY, such as groups and roles, members and their certificates, as well as member groups and roles affiliation and inserts it directly into VOMS database. The users with inactive (expired) Fermilab account are deleted from VOMS service automatically.

6 Integration with SNOW

FERRY integration with SNOW improved user experience by allowing the creation of comprehensive request forms with drop down menus with valid data choices that are populated from FERRY. We have streamlined the request management that allows us to demand approval/rejection from a specific group of managers before processing further requests. The major breakthrough was achieved by implementing a push request to FERRY by using the SNOW Orchestration mechanism that is shown in Figure 3.

	Name	Sequence	Activity Type	Status	Activity Disposition	Result	Time Started	Time Ended	Call Item WF Activity	Updated by	Updated
<input type="checkbox"/>	Empty user	1	Orchestration	Completed	Run	Completed	2018-11-10 16:28:18	2018-11-10 16:28:18	Empty user	keymer	2018-11-10 16:28:18
<input type="checkbox"/>	Successful completion notification	9	Notification	Pending	Run	(empty)	(empty)	(empty)	Successful completion notification	keymer	2018-11-10 16:28:17
<input type="checkbox"/>	Check with FERRY if user is member of th...	3	Orchestration	Completed	Run	CONTINUE. User is a member of group	2018-11-10 16:28:21	2018-11-10 16:28:24	Check with FERRY if user is member of th...	keymer	2018-11-10 16:28:24
<input type="checkbox"/>	Approval by Virtual Organization Approver	5	Approval	Completed	Run		18-11-10 28:25	2018-11-11 05:40:11	Approval by Virtual Organization Approver	guest	2018-11-11 05:40:11
<input type="checkbox"/>	Call FERRY to Add/Remove supersuser	7	Orchestration	Completed	Run	setsupersuser request completed successf...	2018-11-11 05:40:12	2018-11-11 05:40:15	Call FERRY to Add/Remove supersuser	guest	2018-11-11 05:40:15
<input type="checkbox"/>	Wait for FERRY	2	Wait	Skipped	Run	FERRY is reachable	2018-11-10 16:28:18	2018-11-10 16:28:21	Wait for FERRY	keymer	2018-11-10 16:28:21
<input type="checkbox"/>	Check with FERRY if already a supersuser	4	Orchestration	Completed	Run	CONTINUE. User is NOT a Supersuser	2018-11-10 16:28:24	2018-11-10 16:28:25	Check with FERRY if already a supersuser	keymer	2018-11-10 16:28:25
<input type="checkbox"/>	RITM for Add/Remove JobSub Group Supersus...	0	RITM	Completed	Run	RITM Activity Processed	2018-11-10 16:28:17	2018-11-10 16:28:17	RITM for Add/Remove JobSub Group Supersus...	keymer	2018-11-10 16:28:17
<input type="checkbox"/>	Wait for FERRY	6	Wait	Skipped	Run	FERRY is reachable	2018-11-11 05:40:11	2018-11-11 05:40:12	Wait for FERRY	guest	2018-11-11 05:40:12
<input type="checkbox"/>	Manual task: Add/Remove JobSub Supersuser	8	Task	Running	Run		2018-11-11 05:40:15	(empty)	Manual task: Add/Remove JobSub Supersuser	guest	2018-11-11 05:40:15

Fig.3. Example of SNOW form workflow with orchestration.

This workflow will allow the automation of user registration and attribute management. For example, if a user submits a SNOW request to become a member of a particular experiment, a SNOW form allows them to select an experiment and Grid role within the experiment (e.g. Production), then based on a workflow gets an approval for a relevant experiment's coordinator and pushes this information to FERRY. Various services that create

passwd/group files on interactive and grid nodes, update VOMS and dCache information will periodically pull this information from FERRY and change above mentioned files and configuration.

7 Deployment

The database used by the FERRY service is deployed in the centrally managed postgres database cluster maintained by the Core Computing Division at Fermilab. There are production and development instances with nightly backups for disaster recovery. The databases are deployed on a redundant cluster to ensure high availability.

The RESTful API for FERRY is implemented in the GO language. A GIT repo has everything needed to deploy the GO based RESTful API. The steps involved are:

- Execute Git clone
- Populate a configuration file with user/password for the db
- Execute docker-compose up

These actions allow us to deploy FERRY service that is responding on a predetermined port and talking to the configured database. The docker-compose config file specifies how to build the GO code, copies in the config files it needs, runs the GO app, and also starts up separate containers that monitor the database pushing stats to our Graphite service and another container that is updating the FERRY database from central Fermilab databases.

8 Monitoring

We are monitoring both components of the FERRY system. The monitoring data is pushed to Fermilab monitoring infrastructure, Landscape[9]. This data are available through the Open Source visualization tools, like Grafana [10] and Kibana[11]. We are using a docker image named wrouesnel/postgres_exporter[12], a PostgreSQL metric exporter for Prometheus, to monitor database activity and report it to a Graphite/Grafana installation used by a large number of other services and users at Fermilab. This is a generic tool which pushes metrics from a Postgres installation and allows us plot many different metrics relating to database activity so we can detect anomalies and understand usage. The database dashboard is shown in Figure 4.

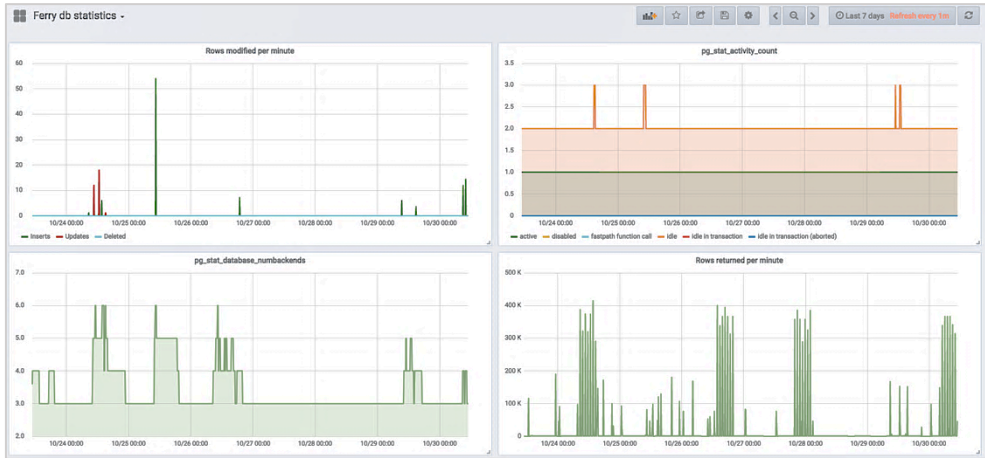


Fig.4. Monitoring database performance using Grafana/Graphite service.

To understand the usage patterns of the individual API calls we are sending FERRY service logs to an ElasticSearch[11] cluster which is also used by many other services and users at Fermilab. The logs of the application are split up into key value pairs which allow us to make plots of anything from individual API invocation counts to plots showing the client locations from where the API was accessed, as well as query content, duration and status. An example of a plot that could be created based on log information is presented in Figure 5.




Fig. 5. Weekly graph of FERRY activities by API requests using Kibana/ES.

9 Conclusion

The FERRY service provides a centralized repository for the access control and job management attributes such as batch and storage access policies, quotas, batch priorities and

NIS attributes for cluster configuration. It maintains a central repository for previously scattered, and sometime obsolete, information related to users' attributes. FERRY provides RESTful APIs that allow uniform data retrieval by numerous services. It enables the use of custom forms from ServiceNow and allows orchestration, so users can request services/changes which will be quickly deployed to services in an automated fashion.

With FERRY deployment we are able to retire unsupported services, such as GUMS and Vulcan and change the service level support from 24x7 basis to 8x5 that reduces the total cost of ownership for the organization.

Overall FERRY reduces operational load on support staff while improving end users experience.

Acknowledgment

This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

References

1. GUMS Retirement: <http://opensciencegrid.org/technology/policy/gums-retire>
2. VULCAN: <https://cmsweb.fnal.gov/bin/view/Software/Vulcanindex>
3. VOMS-ADMIN Retirement: <http://opensciencegrid.org/technology/policy/voms-admin-retire/>
4. VOMS: <http://repository.egi.eu/2012/07/10/voms-2-0-8/>
5. ServiceNow (SNOW): <https://www.servicenow.com/>
6. SNOW Orchestration: <https://www.servicenow.com/products/orchestration.html>
7. HTCondor Overview: <http://opensciencegrid.org/docs/compute-element/htcondor-ce-overview/>
8. <https://github.com/tskirvin/json-passwd>
9. Landscape: https://indico.cern.ch/event/721026/contributions/2964078/attachments/1629322/2615771/Landscape_CERN_201804.pdf
10. Grafana/Graphite: <http://docs.grafana.org/features/datasources/graphite/>
11. ELK Stack: <https://www.elastic.co/elk-stack>
12. https://github.com/wrouesnel/postgres_exporter